

Atlas LSH Neural Networks: Hierarchical Geometry Rather Than Hierarchical Features

One way to interpret a conventional deep neural network is that successive layers learn increasingly abstract features.

Atlas LSH neural networks suggest a different viewpoint.

Rather than building a hierarchy of features, they build a **hierarchy of geometric partitions**.

A multiscale geometric atlas

Each layer divides the input space into progressively finer geometric regions.

- **Layer 1:** Broad geometric patches covering large regions of the input space.
- **Layer 2:** Finer subdivisions within those broad patches.
- **Layer 3:** Further geometric refinement.
- **Deeper layers:** Increasingly localized neighbourhoods.

The network therefore behaves less like a single monolithic function and more like a **multiscale atlas**, where each layer increases the geometric resolution of the representation.

Balanced utilization at every level

Because the routing is determined by LSH bits that are individually approximately 50:50, parameter usage remains naturally balanced throughout the hierarchy.

This has several consequences.

- Every local model receives approximately equal long-run use.
- Training effort is spread throughout the parameter space.
- Model capacity is efficiently utilized rather than concentrating on a small subset of parameters.

Unlike learned routing schemes, this balanced utilization arises from the geometric construction itself.

Geometry-induced co-activation

Balanced utilization does not imply independent parameter usage.

Correlations between LSH bits create **geometry-induced co-activation**:

- nearby inputs tend to activate overlapping collections of local models,
- neighbouring geometric regions share statistical structure,
- while individual parameters still receive approximately equal long-run utilization.

An intuitive description is **fair spread clustering**: computation is distributed evenly across the model while preserving geometric locality.

A different view of depth

This leads to an alternative interpretation of deep learning.

Instead of asking each layer to discover more abstract features, Atlas LSH asks each layer to provide a finer geometric description of the input space.

Depth becomes a process of **progressive geometric refinement**. The model starts with coarse geometric charts and incrementally refines them into increasingly local linear approximations, forming a hierarchical atlas over the input space.

Viewed this way, Atlas LSH is not simply another sparse neural network architecture. It is a geometric computing framework in which balanced utilization and geometry-induced co-activation emerge naturally from the structure of the routing itself.

From co-activation to hierarchical mappings

Geometry-induced co-activation influences not only which parameters are active, but also how the computation itself is assembled.

Each active local model contributes a linear mapping, and these mappings concatenate through successive layers of the network.

The resulting computation has a natural geometric interpretation.

- The first layer applies a mapping appropriate for a broad geometric region of the input space.
- The second layer refines that mapping within a smaller subregion.
- Each subsequent layer provides a progressively more localized refinement.
- The complete network therefore composes a sequence of local linear maps, each conditioned on an increasingly fine geometric partition.

This suggests that Atlas LSH performs computation as a hierarchy of geometric refinements rather than as a sequence of globally shared transformations.

Viewed in this way, the co-activation patterns define compatible chains of local linear maps. Broad geometric behaviour is established early in the network, while deeper layers successively specialize the computation to the increasingly localized neighbourhood occupied by the input.